EAST Search History

Ref	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1 .	101	(creat\$3 or generat\$3) with (cardinalit\$3 or cost\$3) with estimat\$3 with (table\$1 or record\$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/31 11:02
L2	16	(707/2).ccls. and L1	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/31 11:02
L3	12	(optimiz\$3 with execut\$3 with quer\$3) and 2	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/31 11:03
L4	7 ·	3 and @ad<"19991222"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/31 11:03
S1	151	(optimiz\$3 with execut\$3 with quer\$3).ab.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF.	2006/10/31 11:03
S2	40	(optimiz\$3 with execut\$3 with quer\$3).ti.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/30 14:36
S3	157	(optimiz\$3 with execut\$3 with quer\$3).clm.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/30 14:36
S4	45	(creat\$3 or generat\$3) with cardinalit\$3 with estimat\$3	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/30 14:41

EAST Search History

	T		· · · · · · · · · · · · · · · · · · ·	· ·	Т	
S5	9	S1 and S4	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/30 14:38
S6	2	S5 and @ad<"19991222"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/31 11:03
S7	2	("6477523").PN.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/30 14:40
S8	1749	(creat\$3 or generat\$3) with (cardinalit\$3 or cost\$3) with estimat\$3	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/30 15:48
59	25	S1 and S8	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR ,	OFF	2006/10/30 14:42
S10	6	S9 and @ad<"19991222"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/30 15:29
S11	2	("6738755").PN.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/30 15:21
S12	14	S3 and S8 and @ad<"19991222"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/30 15:48

EAST Search History

S13	9	S1 and (707/3).ccls. and @ad<"19991222"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR .	OFF	2006/10/30 15:48
S14	101	(creat\$3 or generat\$3) with (cardinalit\$3 or cost\$3) with estimat\$3 with (table\$1 or record\$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/30 15:49
S15	. 2	S1 and S14	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/10/31 11:02



Subscribe (Full Service) Register (Limited Service, Free) Login

The ACM Digital Library
O The Guide

optimize execution plan query estimates overlap "automatic su

- Maria Land

Feedback Report a problem Satisfaction

Terms used

optimize execution plan query estimates overlap automatic summary table

Found 41,303 of 186,958

Sort results bν

relevance Ţ.

Save results to a Binder Search Tips

Try an Advanced Search Try this search in The ACM Guide

Display results

expanded form

Open results in a new window

Results 1 - 20 of 200 Best 200 shown

Result page: 1 2 3 4 5 6 7 8 9 10

Relevance scale .

SEARCH

Query optimization in distributed networks of autonomous database systems

Fragkiskos Pentaris, Yannis Ioannidis June 2006 ACM Transactions on Database Systems (TODS), Volume 31 Issue 2

Publisher: ACM Press

Full text available: pdf(1.55 MB) Additional Information: full citation, abstract, references, index terms

Large-scale distributed environments, where each node is completely autonomous and

offers services to its peers through external communication, pose significant challenges to query processing and optimization. Autonomy is the main source of the problem, as it results in lack of knowledge about any particular node with respect to the information it can produce and its characteristics, for example, cost of production or quality of produced results. In this article, inspired by e-commerce technolog ...

Keywords: Query optimization

2 Query processing for relational data: Query evaluation using overlapping views:

completeness and efficiency

Gang Gou, Maxim Kormilitsin, Rada Chirkova

June 2006 Proceedings of the 2006 ACM SIGMOD international conference on Management of data SIGMOD '06

Publisher: ACM Press

Full text available: pdf(383.41 KB) Additional Information: full citation, abstract, references, index terms

We study the problem of finding efficient equivalent view-based rewritings of relational queries, focusing on query optimization using materialized views under the assumption that base relations cannot contain duplicate tuples. A lot of work in the literature addresses the problems of answering queries using views and query optimization. However, most of it proposes solutions for special cases, such as for conjunctive queries (CQs) or for aggregate queries only. In addition, most of it addresses ...

Keywords: materialized views, query optimization, rewriting queries using views

A case for dynamic view management Yannis Kotidis, Nick Roussopoulos





December 2001 ACM Transactions on Database Systems (TODS), Volume 26 Issue 4 Publisher: ACM Press

Full text available: pdf(892.57 KB)

Additional Information: full citation, abstract, references, citings, index terms, review

Materialized aggregate views represent a set of redundant entities in a data warehouse that are frequently used to accelerate On-Line Analytical Processing (OLAP). Due to the complex structure of the data warehouse and the different profiles of the users who submit queries, there is need for tools that will automate and ease the view selection and management processes. In this article we present DynaMat, a system that manages dynamic collections of materialized aggregate views in a data warehous ...

Keywords: Data cube, OLAP, data warehousing, materialized views

4 Optimization of dynamic query evaluation plans

Richard L. Cole, Goetz Graefe

May 1994 ACM SIGMOD Record, Proceedings of the 1994 ACM SIGMOD international conference on Management of data SIGMOD '94, Volume 23 Issue 2

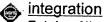
Publisher: ACM Press

Full text available: pdf(1.45 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> <u>terms</u>

Traditional query optimizers assume accurate knowledge of run-time parameters such as selectivities and resource availability during plan optimization, i.e., at compile time. In reality, however, this assumption is often not justified. Therefore, the "static" plans produced by traditional optimizers may not be optimal for many of their actual run-time invocations. Instead, we propose a novel optimization model that assigns the bulk of the optimization effort to compile-time and ...

5 Query Optimization: Joint optimization of cost and coverage of query plans in data



Zaiqing Nie, Subbarao Kambhampati

October 2001 Proceedings of the tenth international conference on Information and knowledge management

Publisher: ACM Press

Full text available: pdf(1.74 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

Existing approaches for optimizing queries in data integration use decoupled strategies-attempting to optimize coverage and cost in two separate phases. Since sources tend to have a variety of access limitations, such phased optimization of cost and coverage can unfortunately lead to expensive planning as well as highly inefficient plans. In this paper we present techniques for joint optimization of cost and coverage of the query plans. Our algorithms search in the space of parallel query plans ...

6 Optimization of query streams using semantic prefetching

Ivan T. Bowman, Kenneth Salem

December 2005 ACM Transactions on Database Systems (TODS), Volume 30 Issue 4

Publisher: ACM Press

Full text available: pdf(1.10 MB)

Additional Information: full citation, abstract, references, index terms

Streams of relational queries submitted by client applications to database servers contain patterns that can be used to predict future requests. We present the Scalpel system, which detects these patterns and optimizes request streams using context-based predictions of future requests. Scalpel uses its predictions to provide a form of semantic prefetching, which involves combining a predicted series of requests into a single request

that can be issued immediately. Scalpel's semantic prefetching ...

Keywords: Prefetching, query streams

⁷ An adaptive query execution system for data integration

Zachary G. Ives, Daniela Florescu, Marc Friedman, Alon Levy, Daniel S. Weld June 1999 ACM SIGMOD Record, Proceedings of the 1999 ACM SIGMOD international

conference on Management of data SIGMOD '99, Volume 28 Issue 2

Publisher: ACM Press

Full text available: pdf(1.59 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

Query processing in data integration occurs over network-bound, autonomous data sources. This requires extensions to traditional optimization and execution techniques for three reasons: there is an absence of quality statistics about the data, data transfer rates are unpredictable and bursty, and slow or unavailable data sources can often be replaced by overlapping or mirrored sources. This paper presents the Tukwila data integration system, designed to support adaptivity at its core using ...

Special issue in parallelism in database systems: Considering data skew factor in multi-way join query optimization for parallel execution

Kien A. Hua, Yo Lung Lo, Honesty C. Young

July 1993 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 2 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(1.43 MB) Additional Information: full citation, abstract, references, citings

A consensus on parallel architecture for very large database management has emerged. This architecture is based on a shared-nothing hardware organization. The computation model is very sensitive to skew in tuple distribution, however. Recently, several parallel join algorithms with dynamic load balancing capabilities have been proposed to address this issue, but none of them consider multi-way join problems. In this article we propose a dynamic load balancing technique for multi-way joins, and i ...

Keywords: load balancing, multi-way join, parallel-database computer, query optimization

9 Industrial session: query processing and optimization: Analyzing plan diagrams of database query optimizers

Naveen Reddy, Jayant R. Haritsa

August 2005 Proceedings of the 31st international conference on Very large data bases VLDB '05

Publisher: VLDB Endowment

Full text available: pdf(362.80 KB) Additional Information: full citation, abstract, references, index terms

A "plan diagram" is a pictorial enumeration of the execution plan choices of a database query optimizer over the relational selectivity space. In this paper, we present and analyze representative plan diagrams on a suite of popular commercial query optimizers for queries based on the TPC-H benchmark. These diagrams, which often appear similar to cubist paintings, provide a variety of interesting insights, including that current optimizers make extremely fine-grained plan choices, which may often ...

10 Join queries with external text sources: execution and optimization techniques

Surajit Chaudhuri, Umeshwar Dayal, Tak W. Yan

May 1995 ACM SIGMOD Record, Proceedings of the 1995 ACM SIGMOD international

conference on Management of data SIGMOD '95, Volume 24 Issue 2

Publisher: ACM Press

Full text available: pdf(1.49 MB)

Additional Information: full citation, abstract, references, citings, index terms

Text is a pervasive information type, and many applications require querying over text sources in addition to structured data. This paper studies the problem of query processing in a system that loosely integrates an extensible database system and a text retrieval system. We focus on a class of conjunctive queries that include joins between text and structured data, in addition to selections over these two types of data. We adapt techniques from distributed query processing and introduce a novel ...

11 Multiple-query optimization

Timos K. Sellis

March 1988 ACM Transactions on Database Systems (TODS), Volume 13 Issue 1

Publisher: ACM Press

Full text available: pdf(2.19 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

Some recently proposed extensions to relational database systems, as well as to deductive database systems, require support for multiple-query processing. For example, in a database system enhanced with inference capabilities, a simple query involving a rule with multiple definitions may expand to more than one actual query that has to be run over the database. It is an interesting problem then to come up with algorithms that process these queries together instead of one query at a time. Th ...

12 Research sessions: query optimization: Robust query processing through

progressive optimization

Volker Markl, Vijayshankar Raman, David Simmen, Guy Lohman, Hamid Pirahesh, Miso Cilimdzic

June 2004 Proceedings of the 2004 ACM SIGMOD international conference on Management of data

Publisher: ACM Press

Full text available: pdf(331.15 KB) Additional Information: full citation, abstract, references

Virtually every commercial query optimizer chooses the best plan for a query using a cost model that relies heavily on accurate cardinality estimation. Cardinality estimation errors can occur due to the use of inaccurate statistics, invalid assumptions about attribute independence, parameter markers, and so on. Cardinality estimation errors may cause the optimizer to choose a sub-optimal plan. We present an approach to query processing that is extremely robust because it is able to detect and re ...

13 Cache investment: integrating query optimization and distributed data placement Donald Kossmann, Michael J. Franklin, Gerhard Drasch, Wig Ag

December 2000 ACM Transactions on Database Systems (TODS), Volume 25 Issue 4

Publisher: ACM Press

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, index Full text available: pdf(210.67 KB) terms

Emerging distributed query-processing systems support flexible execution strategies in which each query can be run using a combination of data shipping and query shipping. As in any distributed environment, these systems can obtain tremendous performance and availability benefits by employing dynamic data caching. When flexible execution and dynamic caching are combined, however, a circular dependency arises: Caching occurs as a by-product of query operator placement, but query operator pl ...

Keywords: cache investment, caching, client-server database systems, data shipping,

dynamic data placement, query optimization, query shipping

14 Research sessions: query processing I: Exploiting statistics on query expressions for



optimization

Nicolas Bruno, Surajit Chaudhuri

June 2002 Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02

Publisher: ACM Press

Full text available: pdf(1.33 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

Statistics play an important role in influencing the plans produced by a query optimizer. Traditionally, optimizers use statistics built over base tables and assume independence between attributes while propagating statistical information through the query plan. This approach can introduce large estimation errors, which may result in the optimizer choosing inefficient execution plans. In this paper, we show how to extend a generic optimizer so that it also exploits statistics built on expression ...

15 Research papers: optimization: Towards a robust query optimizer: a principled and



practical approach

Brian Babcock, Surajit Chaudhuri

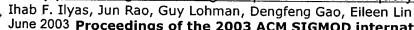
June 2005 Proceedings of the 2005 ACM SIGMOD international conference on Management of data

Publisher: ACM Press

Full text available: pdf(551.42 KB) Additional Information: full citation, abstract, references

Research on query optimization has focused almost exclusively on reducing query execution time, while important qualities such as consistency and predictability have largely been ignored, even though most database users consider these qualities to be at least as important as raw performance. In this paper, we explore how the query optimization process can be made more robust, focusing on the important subproblem of cardinality estimation. The robust cardinality estimation technique that we propo ...

Query processing: Estimating compilation time of a query optimizer



June 2003 Proceedings of the 2003 ACM SIGMOD international conference on Management of data

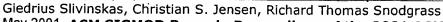
Publisher: ACM Press

Full text available: pdf(292.76 KB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> <u>terms</u>

A query optimizer compares alternative plans in its search space to find the best plan for a given query. Depending on the search space and the enumeration algorithm, optimizers vary in their compilation time and the quality of the execution plan they can generate. This paper describes a compilation time estimator that provides a quantified estimate of the optimizer compilation time for a given query. Such an estimator is useful for automatically choosing the right level of optimization in comme ...

17 Adaptable query optimization and evaluation in temporal middleware



May 2001 ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international conference on Management of data SIGMOD '01, Volume 30 Issue 2

Publisher: ACM Press

Full text available: pdf(232.65 KB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u>

Time-referenced data are pervasive in most real-world databases. Recent advances in temporal query languages show that such database applications may benefit substantially from built-in temporal support in the DBMS. To achieve this, temporal query optimization and evaluation mechanisms must be provided, either within the DBMS proper or as a source level translation from temporal queries to conventional SQL. This paper proposes a new approach: using a middleware component on top of a conventio ...

18 Research sessions: non-standard query processing: Optimization of query streams



<u>ا</u> 🕲

using semantic prefetching

Ivan T. Bowman, Kenneth Salem

June 2004 Proceedings of the 2004 ACM SIGMOD international conference on Management of data

Publisher: ACM Press

Full text available: pdf(224.40 KB) Additional Information: full citation, abstract, references

Streams of relational queries submitted by client applications to database servers contain patterns that can be used to predict future requests. We present the Scalpel system, which detects these patterns and optimizes request streams using context-based predictions of future requests. Scalpel uses its predictions to provide a form of semantic prefetching, which involves combining a predicted series of requests into a single request that can be issued immediately. Scalpel's semantic prefetching ...

19 Query processing and optimization in Oracle Rdb

Gennady Antoshenkov, Mohamed Ziauddin

December 1996 The VLDB Journal — The International Journal on Very Large Data
Bases, Volume 5 Issue 4

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(92.62 KB) Additional Information: full citation, abstract, citings, index terms

This paper contains an overview of the technology used in the query processing and optimization component of Oracle Rdb, a relational database management system originally developed by Digital Equipment Corporation and now under development by Oracle Corporation. Oracle Rdb is a production system that supports the most demanding database applications, runs on multiple platforms and in a variety of environments.

Keywords: Dynamic optimization, Optimizer, Query transformation, Relational database, Sampling

20 Multiway spatial joins



Nikos Mamoulis, Dimitris Papadias

December 2001 ACM Transactions on Database Systems (TODS), Volume 26 Issue 4

Publisher: ACM Press

Full text available: pdf(2.04 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms, review

Due to the evolution of Geographical Information Systems, large collections of spatial data having various thematic contents are currently available. As a result, the interest of users is not limited to simple spatial selections and joins, but complex query types that implicate numerous spatial inputs become more common. Although several algorithms have been proposed for computing the result of pairwise spatial joins, limited work exists on processing and optimization of *multiway spatial join* ...

Keywords: Multiway joins, query processing, spatial joins

Results (page 1): optimize execution plan query estimates overlap "automatic summary ta... Page 7 of 7

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10 next

ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

<u>Terms of Usage Privacy Policy Code of Ethics Contact Us</u>

Useful downloads: Adobe Acrobat Q QuickTime Windows Media Player Real Player



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: • The ACM Digital Library C The Guide

optimize execution plan query estimates overlap materialized \

SEARCH



Feedback Report a problem Satisfaction

Terms used optimize execution plan query estimates overlap materialized view

Found 61,382 of 186,958

Sort results by

relevance ∇

Save results to a Binder Search Tips

Try an Advanced Search Try this search in The ACM Guide

Display results

expanded form

Open results in a new window

Result page: 1 2 3 4 5 6 7 8 9 10

next

Results 1 - 20 of 200 Best 200 shown

Relevance scale

Materialized view selection and maintenance using multi-query optimization

Hoshi Mistry, Prasan Roy, S. Sudarshan, Krithi Ramamritham

May 2001 ACM SIGMOD Record, Proceedings of the 2001 ACM SIGMOD international conference on Management of data SIGMOD '01, Volume 30 Issue 2

Publisher: ACM Press

Full text available: pdf(199.46 KB)

Additional Information: full citation, abstract, references, citings, index terms

Materialized views have been found to be very effective at speeding up queries, and are increasingly being supported by commercial databases and data warehouse systems. However, whereas the amount of data entering a warehouse and the number of materialized views are rapidly increasing, the time window available for maintaining materialized views is shrinking. These trends necessitate efficient techniques for the maintenance of materialized views.

In this paper, we show how to find an ...

² Cost-based optimization of decision support queries using transient-views



Subbu N. Subramanian, Shivakumar Venkataraman

June 1998 ACM SIGMOD Record, Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98, Volume 27 Issue 2

Publisher: ACM Press

Full text available: pdf(1.58 MB)

Additional Information: full citation, abstract, references, citings, index terms

Next generation decision support applications, besides being capable of processing huge amounts of data, require the ability to integrate and reason over data from multiple, heterogeneous data sources. Often, these data sources differ in a variety of aspects such as their data models, the query languages they support, and their network protocols. Also, typically they are spread over a wide geographical area. The cost of processing decision support queries in such a setting is quite high. Ho ...

3 Efficient and extensible algorithms for multi query optimization



Prasan Roy, S. Seshadri, S. Sudarshan, Siddhesh Bhobe May 2000 ACM SIGMOD Record, Proceedings of the 2000 ACM SIGMOD international conference on Management of data SIGMOD '00, Volume 29 Issue 2

Publisher: ACM Press

Full text available: pdf(197.59 KB) Additional Information: full citation, abstract, references, citings, index terms

Complex queries are becoming commonplace, with the growing use of decision support systems. These complex queries often have a lot of common sub-expressions, either within a single query, or across multiple such queries run as a batch. Multiquery optimization aims at exploiting common sub-expressions to reduce evaluation cost. Multiquery optimization has hither-to been viewed as impractical, since earlier algorithms were exhaustive, and explore a doubly exponential search space.

In t ...

4 Query processing for relational data: Query evaluation using overlapping views:

completeness and efficiency

Gang Gou, Maxim Kormilitsin, Rada Chirkova

June 2006 Proceedings of the 2006 ACM SIGMOD international conference on Management of data SIGMOD '06

Publisher: ACM Press

Full text available: pdf(383.41 KB) Additional Information: full citation, abstract, references, index terms

We study the problem of finding efficient equivalent view-based rewritings of relational queries, focusing on query optimization using materialized views under the assumption that base relations cannot contain duplicate tuples. A lot of work in the literature addresses the problems of answering queries using views and query optimization. However, most of it proposes solutions for special cases, such as for conjunctive queries (CQs) or for aggregate queries only. In addition, most of it addresses ...

Keywords: materialized views, query optimization, rewriting queries using views

5 Research sessions: query processing I: Exploiting statistics on query expressions for



optimization

Nicolas Bruno, Surajit Chaudhuri

June 2002 Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02

Publisher: ACM Press

Full text available: pdf(1.33 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> <u>terms</u>

Statistics play an important role in influencing the plans produced by a query optimizer. Traditionally, optimizers use statistics built over base tables and assume independence between attributes while propagating statistical information through the query plan. This approach can introduce large estimation errors, which may result in the optimizer choosing inefficient execution plans. In this paper, we show how to extend a generic optimizer so that it also exploits statistics built on expression ...

6 Query evaluation techniques for large databases

Goetz Graefe

June 1993 ACM Computing Surveys (CSUR), Volume 25 Issue 2

Publisher: ACM Press

Full text available: pdf(9.37 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms, review

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models

exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

7 The state of the art in distributed query processing

Donald Kossmann

December 2000 ACM Computing Surveys (CSUR), Volume 32 Issue 4

Publisher: ACM Press

Full text available: pdf(455.39 KB)

Additional Information: full citation, abstract, references, citings, index terms

Distributed data processing is becoming a reality. Businesses want to do it for many reasons, and they often must do it in order to stay competitive. While much of the infrastructure for distributed data processing is already there (e.g., modern network technology), a number of issues make distributed data processing still a complex undertaking: (1) distributed systems can become very large, involving thousands of heterogeneous sites including PCs and mainframe server machines; (2) the stat ...

Keywords: caching, client-server databases, database application systems, dissemination-based information systems, economic models for query processing, middleware, multitier architectures, query execution, query optimization, replication, wrappers

8 Query optimization in distributed networks of autonomous database systems

Fragkiskos Pentaris, Yannis Ioannidis

June 2006 ACM Transactions on Database Systems (TODS), Volume 31 Issue 2

Publisher: ACM Press

Full text available: pdf(1.55 MB) Additional Information: full citation, abstract, references, index terms

Large-scale distributed environments, where each node is completely autonomous and offers services to its peers through external communication, pose significant challenges to query processing and optimization. Autonomy is the main source of the problem, as it results in lack of knowledge about any particular node with respect to the information it can produce and its characteristics, for example, cost of production or quality of produced results. In this article, inspired by e-commerce technolog ...

Keywords: Query optimization

⁹ Research papers: OLAP: Efficient computation of multiple group by queries

Zhimin Chen, Vivek Narasayya

June 2005 Proceedings of the 2005 ACM SIGMOD international conference on Management of data

Publisher: ACM Press

Full text available: pdf(371.92 KB) Additional Information: full citation, abstract, references

Data analysts need to understand the quality of data in the warehouse. This is often done by issuing many Group By queries on the sets of columns of interest. Since the volume of data in these warehouses can be large, and tables in a data warehouse often contain many columns, this analysis typically requires executing a large number of Group By queries, which can be expensive. We show that the performance of today's database

systems for such data analysis is inadequate. We also show that the pro ...

10 Query processing: Estimating compilation time of a query optimizer

Ihab F. Ilyas, Jun Rao, Guy Lohman, Dengfeng Gao, Eileen Lin

June 2003 Proceedings of the 2003 ACM SIGMOD international conference on Management of data

Publisher: ACM Press

Full text available: pdf(292.76 KB)

Additional Information: full citation, abstract, references, citings, index terms

A query optimizer compares alternative plans in its search space to find the best plan for a given query. Depending on the search space and the enumeration algorithm, optimizers vary in their compilation time and the quality of the execution plan they can generate. This paper describes a compilation time estimator that provides a quantified estimate of the optimizer compilation time for a given query. Such an estimator is useful for automatically choosing the right level of optimization in comme ...

Answering queries using views: A survey

Alon Y. Halevv

December 2001 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 10 Issue 4

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(308.74 KB) Additional Information: full citation, abstract, citings, index terms

The problem of answering queries using views is to find efficient methods of answering a query using a set of previously defined materialized views over the database, rather than accessing the database relations. The problem has recently received significant attention because of its relevance to a wide variety of data management problems. In query optimization, finding a rewriting of a query using a set of materialized views can yield a more efficient query execution plan. To support the separat ...

Keywords: Data integration, Date warehousing, Materialized views, Query optimization, Survey, Web-site management

12 A case for dynamic view management

Yannis Kotidis, Nick Roussopoulos

December 2001 ACM Transactions on Database Systems (TODS), Volume 26 Issue 4

Publisher: ACM Press

Full text available: pdf(892.57 KB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms, review

Materialized aggregate views represent a set of redundant entities in a data warehouse that are frequently used to accelerate On-Line Analytical Processing (OLAP). Due to the complex structure of the data warehouse and the different profiles of the users who submit queries, there is need for tools that will automate and ease the view selection and management processes. In this article we present DynaMat, a system that manages dynamic collections of materialized aggregate views in a data warehous ...

Keywords: Data cube, OLAP, data warehousing, materialized views

13 Optimization of query streams using semantic prefetching

Ivan T. Bowman, Kenneth Salem

December 2005 ACM Transactions on Database Systems (TODS), Volume 30 Issue 4

Publisher: ACM Press

Full text available: pdf(1.10 MB)

Additional Information: full citation, abstract, references, index terms

Streams of relational queries submitted by client applications to database servers contain patterns that can be used to predict future requests. We present the Scalpel system, which detects these patterns and optimizes request streams using context-based predictions of future requests. Scalpel uses its predictions to provide a form of semantic prefetching, which involves combining a predicted series of requests into a single request that can be issued immediately. Scalpel's semantic prefetching ...

Keywords: Prefetching, query streams

14 Cache investment: integrating query optimization and distributed data placement

Donald Kossmann, Michael J. Franklin, Gerhard Drasch, Wig Ag
December 2000 ACM Transactions on Database Systems (TODS), Volume 25 Issue 4

Publisher: ACM Press

Full text available: pdf(210.67 KB)

Additional Information: full citation, abstract, references, citings, index terms

Emerging distributed query-processing systems support flexible execution strategies in which each query can be run using a combination of data shipping and query shipping. As in any distributed environment, these systems can obtain tremendous performance and availability benefits by employing dynamic data caching. When flexible execution and dynamic caching are combined, however, a circular dependency arises: Caching occurs as a by-product of query operator placement, but query operator pl ...

Keywords: cache investment, caching, client-server database systems, data shipping, dynamic data placement, query optimization, query shipping

15 Query processing for relational data: Supporting ad-hoc ranking aggregates

Chengkai Li, Kevin Chen-Chuan Chang, Ihab F. Ilyas

June 2006 Proceedings of the 2006 ACM SIGMOD international conference on Management of data SIGMOD '06

Publisher: ACM Press

Full text available: pdf(344.23 KB) Additional Information: full citation, abstract, references, index terms

This paper presents a principled framework for efficient processing of ad-hoc top-k (ranking) aggregate queries, which provide the k groups with the highest aggregates as results. Essential support of such queries is lacking in current systems, which process the queries in a naïve materialize-group-sort scheme that can be prohibitively inefficient. Our framework is based on three fundamental principles. The Upper-Bound Principle dictates the requirements of early pruning, and ...

Keywords: OLAP, aggregate query, decision support, ranking, top-k query processing

16 Cost-based query scrambling for initial delays

Tolga Urhan, Michael J. Franklin, Laurent Amsaleg

June 1998 ACM SIGMOD Record, Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98, Volume 27 Issue 2

Publisher: ACM Press

Full text available: pdf(1.81 MB)

Additional Information: full citation, abstract, references, citings, index terms

Remote data access from disparate sources across a wide-area network such as the Internet is problematic due to the unpredictable nature of the communications medium and the lack of knowledge about the load and potential delays at remote sites. Traditional,

static, query processing approaches break down in this environment because they are unable to adapt in response to unexpected delays. Query scrambling has been proposed to address this problem. Scrambling modifies query execution plans o ...

17 Research sessions: data integration: Adapting to source properties in processing



data integration queries

Zachary G. Ives, Alon Y. Halevy, Daniel S. Weld

June 2004 Proceedings of the 2004 ACM SIGMOD international conference on Management of data

Publisher: ACM Press

Full text available: pdf(197.27 KB) Additional Information: full citation, abstract, references

An effective query optimizer finds a query plan that exploits the characteristics of the source data. In data integration, little is known in advance about sources' properties, which necessitates the use of adaptive query processing techniques to adjust query processing on-the-fly. Prior work in adaptive query processing has focused on compensating for delays and adjusting for mis-estimated cardinality or selectivity values. In this paper, we present a generalized architecture for adaptiv ...

18 Query execution techniques for caching expensive methods



Joseph M. Hellerstein, Jeffrey F. Naughton

June 1996 ACM SIGMOD Record, Proceedings of the 1996 ACM SIGMOD international conference on Management of data SIGMOD '96, Volume 25 Issue 2

Publisher: ACM Press

Full text available: pdf(1.53 MB)

Additional Information: full citation, abstract, references, citings, index terms

Object-Relational and Object-Oriented DBMSs allow users to invoke time-consuming ("expensive") methods in their queries. When queries containing these expensive methods are run on data with duplicate values, time is wasted redundantly computing methods on the same value. This problem has been studied in the context of programming languages, where "memoization" is the standard solution. In the database literature, sorting has been proposed to deal with this problem. We compare these approaches al ...

19 Research sessions: non-standard query processing: Optimization of query streams



using semantic prefetching

Ivan T. Bowman, Kenneth Salem

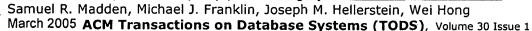
June 2004 Proceedings of the 2004 ACM SIGMOD international conference on Management of data

Publisher: ACM Press

Full text available: pdf(224.40 KB) Additional Information: full citation, abstract, references

Streams of relational queries submitted by client applications to database servers contain patterns that can be used to predict future requests. We present the Scalpel system, which detects these patterns and optimizes request streams using context-based predictions of future requests. Scalpel uses its predictions to provide a form of semantic prefetching, which involves combining a predicted series of requests into a single request that can be issued immediately. Scalpel's semantic prefetching ...

²⁰ TinyDB: an acquisitional query processing system for sensor networks



Publisher: ACM Press

Full text available: pdf(1.67 MB) Additional Information: full citation, abstract, references, index terms

We discuss the design of an acquisitional query processor for data collection in sensor

networks. Acquisitional issues are those that pertain to where, when, and how often data is physically acquired (sampled) and delivered to query processing operators. By focusing on the locations and costs of acquiring data, we are able to significantly reduce power consumption over traditional passive systems that assume the a priori existence of data. We discuss simple extensions to SQL for controlli ...

Keywords: Query processing, data acquisition, sensor networks

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

<u>Terms of Usage Privacy Policy Code of Ethics Contact Us</u>

Useful downloads: Adobe Acrobat Q QuickTime Windows Media Player Real Player



Home | Login | Logout | Access Information | Alerts |

Welcome United States Patent and Trademark Office

EDESearch Results

BROWSE

SEARCH

IEEE XPLORE GUIDE

Results for "((optimize execution plan query estimates overlap statistic materialized view)<in>metadata)" Your search matched 0 documents.

☑ e-mail

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» Search Options

View Session History

Modify Search

((optimize execution plan query estimates overlap statistic materialized view)<in>met

» Key

IEEE Journal or Magazine

IEEE JNL IEE JNL

New Search

IEE Journal or Magazine

IEEE CNF

IEEE Conference

Proceeding

IEE CNF

IEE Conference

Proceeding

IEEE STD IEEE Standard

Check to search only within this results set

Display Format: © Citation © Citation & Abstract

No results were found.

Please edit your search criteria and try again. Refer to the Help pages if you need assistan

search.

Indexed by inspec' Help Contact Us Privacy &:

© Copyright 2006 IEEE -

Sign in



Images Video News <u>Maps</u>

optimize execution plan query estimates "auto Search



Advanced Search

Web Results 1 - 2 of 2 for optimize execution plan query estimates "automatic summary table" overlap st

Tip: Try removing quotes from your search to get more results.

Query optimization technique for obtaining improved cardinality ... A technique for optimizing execution of a query that accesses data stored on ... associated with an automatic summary table, a materialized view or a view. ... www.freepatentsonline.com/20040181521.html - 82k - Cached - Similar pages

[PDF] Architecting Portal Solutions al Solutions

File Format: PDF/Adobe Acrobat optimizes the query, developing an execution plan in which the query has ... allows administrators to define materialized views of data in a set of ... www.redbooks.ibm.com/redbooks/pdfs/sg247011.pdf - Similar pages

Free! Speed up the web. Download the Google Web Accelerator.

optimize execution plan query estimated Search

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

©2006 Google

Sign in



News Maps more » Images Video

optimize cardinality execution plan query estimates Search

Advanced Search

Web Results 1 - 10 of about 636 for optimize cardinality execution plan query estimates overlap statistic n

[PDF] Automatic Management of Statistics on Query Expressions in ...

File Format: PDF/Adobe Acrobat - View as HTML

query execution plans, query optimizers estimate the cost of each ... Optimizing queries using materialized views: A practical, scalable solution. ...

research.microsoft.com/users/nicolasb/papers/proposal.pdf - Similar pages

Query optimization technique for obtaining improved cardinality ...

In particular, improved cardinality estimates are generated for one or more query execution plans for the query using statistics of one or more pre-defined ... www.freepatentsonline.com/20040181521.html - 82k - Cached - Similar pages

Optimizing SQL Statements

Estimate table cardinality for tables without statistics or for tables whose ... Optimization

(determining the execution plan) takes place before the ...

www.lc.leidenuniv.nl/awcourse/oracle/server.920/a96533/sql_1016.htm - 82k -

Cached - Similar pages

[PDF] Efficient Creation of Statistics over Query Expressions

File Format: PDF/Adobe Acrobat - View as HTML

fore, quality of the query execution plan depends on the ac- curacy of cost estimates.

Cost estimates, in turn, crucially, depend on cardinality ...

www.cs.brown.edu/courses/cs227/Papers/AutoAdmin/buildsits.pdf - Similar pages

[РРТ] Using State Modules for Adaptive Query Processing

File Format: Microsoft Powerpoint - View as HTML

Eddy executes different query execution plans for different parts of data ... To guard

against wrong cardinality estimates; To simultaneously optimize for ...

www.cs.umd.edu/~amol/talks/Tutorial-COMAD.ppt - Similar pages

[PDF] Rank-aware Query Optimization

File Format: PDF/Adobe Acrobat - View as HTML

materialized views or special indexes to enhance the pere-. sponse time of top-k queries

are ... space to generate rank-aware query execution plans. Rank- ...

www.cs.uwaterloo.ca/~ilyas/papers/ilyassigmod04.pdf - Similar pages

грет Information Integration

File Format: Microsoft Powerpoint - View as HTML

No central administration, Uncontrolled source content overlap, Lack of source statistics.

Tradeoffs between query plan cost, coverage, quality etc. ...

rakaposhi.eas.asu.edu/cse494/notes/f05-data-integration.ppt - Similar pages

[PS] Exploiting Constraint>Like Data Characterizations in Query ...

File Format: Adobe PostScript - View as Text

be used in query optimization and execution. ... ing it dicult to estimate cardinality

because of statistical. correlation between values in the columns. ...

www.cse.yorku.ca/~jarek/papers/sigmod01/paper.ps - Similar pages

[PDF] 1 Query Selectivity Estimation via Data Mining

File Format: PDF/Adobe Acrobat - View as HTML

quality execution plans for complex join queries. We are thus motivated to, propose

building new statistics over non base-relations for better estimates ... www.cse.yorku.ca/~jarek/papers/iis04/paper.pdf - Similar pages

[PDF] Query Processing and Optimization on the Web

File Format: PDF/Adobe Acrobat - View as HTML execution plans and select the best one. Traditional query optimization ... cardinality statistics in a database and provides cost estimates to the rule ... www.cs.purdue.edu/homes/mourad/publications/dapd2004.pdf - Similar pages

> Result Page: 1 2 3 4 5 6 7 8 9 10

<u>Next</u>

Free! Speed up the web. Download the Google Web Accelerator.

optimize cardinality execution plan g Search

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

©2006 Google